

A Complexity-Regularized Quantization Approach to Nonlinear Dimensionality Reduction

Maxim Raginsky

Beckman Institute and the University of Illinois

405 N Mathews Ave, Urbana, IL 61801, USA

Email: maxim@uiuc.edu

Abstract—We consider the problem of nonlinear dimensionality reduction: given a training set of high-dimensional data whose “intrinsic” low dimension is assumed known, find a feature extraction map to low-dimensional space, a reconstruction map back to high-dimensional space, and a geometric description of the dimension-reduced data as a smooth manifold. We introduce a complexity-regularized quantization approach for fitting a Gaussian mixture model to the training set via a Lloyd algorithm. Complexity regularization controls the trade-off between adaptation to the local shape of the underlying manifold and global geometric consistency. The resulting mixture model is used to design the feature extraction and reconstruction maps and to define a Riemannian metric on the low-dimensional data. We also sketch a proof of consistency of our scheme for the purposes of estimating the unknown underlying pdf of high-dimensional data.

I. INTRODUCTION

When dealing with high volumes of vector-valued data of some large dimension n , it is often assumed that the data possess some intrinsic geometric description in a space of unknown dimension $k < n$ and that the high dimensionality arises from an unknown stochastic mapping of \mathbb{R}^k into \mathbb{R}^n . We can pose the problem of *nonlinear dimensionality reduction* (NLDR) [1], [2] as follows: given raw data with values in \mathbb{R}^n , we wish to obtain optimal estimates of the intrinsic dimension k and of the stochastic map with the purpose of modeling the intrinsic geometry of the data in \mathbb{R}^k .

One typically considers the following set-up: we are given a sample $X^N \equiv (X_1, \dots, X_N)$, where X_i are i.i.d. according to an unknown absolutely continuous distribution P^* . The corresponding pdf f^* has to be estimated from the observation as $\hat{f}_N \equiv \hat{f}_N(X^N)$. The intrinsic dimension k of the data may not be known in advance and would also have to be estimated as $\hat{k}_N \equiv \hat{k}_N(X^N)$. Since the pdf f^* is assumed to arise from a stochastic map of the low-dimensional space \mathbb{R}^k into the high-dimensional space \mathbb{R}^n , we can use our knowledge about k and f^* in order to make inferences about the intrinsic geometry of the data. In the absence of such knowledge, any such inference has to be made based on the estimates \hat{k}_N and \hat{f}_N . In this paper we introduce a complexity-regularized quantization approach to NLDR, assuming that the intrinsic dimension k of the data is given (e.g., as a maximum-likelihood estimate [3]).

II. SMOOTH MANIFOLDS AND THEIR NOISY EMBEDDINGS

We begin with a quick sketch of some notions about smooth manifolds [4]. A *smooth manifold* of dimension k is a set M together with a collection $\mathcal{A} = \{(U_l, \varphi_l) : l \in \Lambda\}$, where the sets $U_l \subset M$ cover M and each map φ_l is a bijection of U_l

onto an open set $\varphi_l(U_l) \subset \mathbb{R}^k$, such that for all l, l' with $U_l \cap U_{l'} \neq \emptyset$ the map $\varphi_{l'} \circ \varphi_l^{-1} : \varphi_l(U_l \cap U_{l'}) \rightarrow \varphi_{l'}(U_l \cap U_{l'})$ is smooth. The pairs (U_l, φ_l) are called *charts* of M , and the entire collection \mathcal{A} is referred to as an *atlas*. Intuitively, the charts describe the points of M by *local* coordinates: given $p \in M$ and a chart $(U_l \ni p, \varphi_l)$, φ_l maps any point q “near p ” (i.e., $q \in U_l$) to an element of $\varphi_l(U_l) \subset \mathbb{R}^k$. Smoothness of the transition maps $\varphi_{l'} \circ \varphi_l^{-1}$ ensures that local coordinates of a point transform differentiably under a change of chart.

Assuming that M is compact, we can always choose the atlas \mathcal{A} in such a way that the indexing set Λ is finite and each $\varphi_l(U_l)$ is an open ball of radius r_l [4, Thm. 3.3] (one can always set $r_l \equiv 1$ for all $l \in \Lambda$, but we choose not to do this for greater flexibility in modeling).

The next notion we need is that of a *tangent space* to M at point p , denoted by $T_p M$. Let $I \subset \mathbb{R}$ be an open interval such that $0 \in I$. Consider the set of all curves $\xi : I \rightarrow M$ such that $\xi(0) = p$. Then for any chart $(U_l \ni p, \varphi_l)$ we have a function $\xi_l \triangleq \varphi_l \circ \xi : I \rightarrow \mathbb{R}^k$, such that $\xi_l(t) \in \varphi_l(U_l)$ for all t in a sufficiently small neighborhood of 0. We say that two such curves ξ, ξ' are equivalent iff $d\xi_{l,j}(t)/dt|_{t=0} = d\xi'_{l,j}(t)/dt|_{t=0}$, $j = 1, \dots, k$, for all $l \in \Lambda$ such that $U_l \ni p$, where $\xi_{l,j}(t)$ are the components of $\xi_l(t)$. The resulting set of equivalence classes has the structure of a vector space of dimension k , and is precisely the tangent space $T_p M$. Intuitively, $T_p M$ allows us to “linearize” M around p . Note that, although all the tangent spaces $T_p M, p \in M$ are isomorphic to each other and to \mathbb{R}^k , there is no meaningful way to add elements of $T_p M$ and $T_q M$ with p, q distinct.

Next, we specify the class of stochastic embeddings dealt with in this paper. Consider three random variables L, Y, X , where L takes values in the finite set Λ with $w_l \triangleq \Pr(L = l)$, Y takes values in \mathbb{R}^k , and X takes values in \mathbb{R}^n . Conditional distributions of Y given L and of X given Y, L are assumed to be absolutely continuous and described by densities $f_{Y|L}$ and $f_{X|Y,L}$, respectively. Since for a compact M the images $\varphi_l(U_l)$ of charts in \mathcal{A} are open balls of radii r_l , let us suppose that the conditional mean $m_l(Y) \equiv \mathbb{E}[Y|L = l]$ is the center of $\varphi_l(U_l)$ [we can therefore take $m_l(Y) = 0$ for all $l \in \Lambda$] and that the largest eigenvalue of the conditional covariance matrix $K_l(Y) \equiv \mathbb{E}[YY^t|L = l]$ of Y given $L = l$ is equal to r_l^2 . It is convenient to think of the eigenvectors $e_1^{(l)}, \dots, e_k^{(l)}$ of $K_l(Y)$ as giving a basis of the tangent space $T_{\varphi_l^{-1}(0)} M$. The unconditional density f_X of X is the finite mixture $f_X(x) = \sum_{l \in \Lambda} w_l f_l(x)$, where $f_l(x) \triangleq \int_{\mathbb{R}^k} f_{X|Y,L}(x|y, l) f_{Y|L}(y|l) dy$.

The resulting pdf follows the local structure of the manifold M and accounts both for low- and high-dimensional noise.

As an example [5], let all $f_{Y|L}(y|l)$ be k -dimensional zero-mean Gaussians with unit covariance matrices, $f_{Y|L}(y|l) = \mathcal{N}(y; 0, I) \equiv (2\pi)^{-k/2} \exp(-\frac{1}{2}y^t y)$, and $f_{X|YL}(x|y, l) = \mathcal{N}(x; \mu_l + A_l y, \Sigma_l)$, $\forall l \in \Lambda$, for some means $\mu_l \in \mathbb{R}^n$, covariance matrices Σ_l , and $n \times k$ matrices A_l , so that $f_X(x) = \sum_{l \in \Lambda} w_l \mathcal{N}(x; \mu_l, A_l A_l^t + \Sigma_l)$.

III. COMPLEXITY-REGULARIZED MIXTURE MODELS

Consider a random vector $X \in \mathbb{R}^n$ with an absolutely continuous distribution P_f , described by a pdf f . We wish to find a mixture model that would not only yield a good “local” approximation to f , but also have low complexity, where the precise notion of complexity depends on application.

In order to set this up quantitatively, we use a complexity-regularized adaptation of the quantizer mismatch approach of Gray and Linder [6]. We seek a finite collection $\Gamma = \{g_m : m \in \mathcal{M}\}$ of pdf’s from a class \mathcal{G} of “admissible” models and a measurable partition $\mathcal{R} = \{R_m : m \in \mathcal{M}\}$ of \mathbb{R}^n that would minimize the objective function

$$\bar{I}_f(\mathcal{R}, \Gamma) \triangleq \sum_{m \in \mathcal{M}} P_f(R_m) [D(f_m \| g_m) + \mu \Phi_\Gamma(g_m)], \quad (1)$$

where f_m is the pdf defined as $1_{\{x \in R_m\}} f(x) / P_f(R_m)$, $D(\cdot \| \cdot)$ is the relative entropy, $\Phi_\Gamma(g_m)$ is a regularization functional that quantifies the complexity of the m th model pdf relative to the entire collection Γ , and $\mu \geq 0$ is the parameter that controls the trade-off between the relative-entropy (mismatch) term and the complexity term.

This minimization problem can be posed as a *complexity-constrained quantization problem* with an encoder $\alpha : \mathbb{R}^n \rightarrow \mathcal{M}$ corresponding to the partition $\mathcal{R} = \{R_m\}$ through $\alpha(x) = m$ if $x \in R_m$, a decoder $\beta : \mathcal{M} \rightarrow \mathcal{G}$ defined by $\beta(m) = g_m$, and a length function $\ell : \mathcal{M} \rightarrow \{0, 1, 2, \dots\}$ satisfying the Kraft inequality $\sum_{m \in \mathcal{M}} e^{-\ell(m)} \leq 1$. In order to describe the encoder and to quantify the performance of the quantization scheme, we need to choose a distortion measure between an input vector and an encoder output in such a way that minimizing average distortion would yield the \bar{I} -functional (1) of the corresponding partition and codebook.

Consider the distortion $\rho(x, m) \triangleq \ln(f(x)/g_m(x)) + \ell(m) + \mu \Phi_\Gamma(g_m)$ (this is not a distortion measure in the strict sense since it can be negative, but its expectation with respect to f is nonnegative by the divergence inequality). For a given codebook Γ and length function ℓ , the optimal encoder is the minimum-distortion encoder $\alpha(x) = \arg \min_{m \in \mathcal{M}} \rho(x, m)$ with ties broken arbitrarily. The resulting partition $\mathcal{R} = \{R_m\}$ yields the average distortion

$$\begin{aligned} E_f \rho(X, \alpha(X)) &= \sum_{m \in \mathcal{M}} p_m \left[\ell(m) + \mu \Phi_\Gamma(g_m) \right. \\ &\quad \left. + \int_{R_m} f_m(x) \ln \frac{p_m f_m(x)}{g_m(x)} dx \right], \end{aligned}$$

where $p_m \triangleq P_f(R_m)$. Then

$$\begin{aligned} E_f \rho(X, \alpha(X)) &= \sum_{m \in \mathcal{M}} p_m \left[D(f_m \| g_m) \right. \\ &\quad \left. + \ln \frac{p_m}{e^{-\ell(m)}} + \mu \Phi_\Gamma(g_m) \right] \\ &\geq \sum_{m \in \mathcal{M}} p_m [D(f_m \| g_m) + \mu \Phi_\Gamma(g_m)], \end{aligned}$$

with equality if and only if $\ell(m) = -\ln p_m$. Thus, the optimal decoder and length function for a given partition are such that the average ρ -distortion is precisely the \bar{I} -functional. We can therefore iterate the optimality properties of the encoder, decoder and length function in a Lloyd-type descent algorithm; this can only decrease average distortion and thus the \bar{I} -functional. Note that the $\ln f(x)$ term in $\rho(x, m)$ does not affect the minimum-distortion encoder. Thus, as far as the encoder is concerned, the distortion measure $\rho_0(x, m) \triangleq -\ln g_m(x) + \ell(m) + \mu \Phi_\Gamma(g_m)$ is equivalent to ρ .

When the distribution of X is unknown, we can take a sufficiently large training sample $X^N = (X_1, \dots, X_N)$ and use a Lloyd descent algorithm to empirically design a mixture model for the data:

- 1) **Initialization:** begin with an initial codebook $\Gamma = \{g_m^{(0)} : m \in \mathcal{M}\} \subset \mathcal{G}$, where \mathcal{G} is the class of admissible models, and a length function $\ell^{(0)} : \mathcal{M} \rightarrow \{0, 1, 2, \dots\}$. Set iteration number $r = 1$, pick a convergence threshold ϵ , and let D_0 be the average ρ_0 -distortion of the initial codebook.
- 2) **Minimum-distortion encoder:** encode each sample X_i into the index $\alpha^{(r)}(X_i) = \arg \min_{m \in \mathcal{M}} \rho_0(X_i, g_m^{(r-1)})$.
- 3) **Centroid decoder:** update the codebook by minimizing over all $g \in \mathcal{G}$ the empirical conditional expectation

$$E[\rho_0(X, g) | \alpha^{(r)}(X) = m] \equiv \frac{1}{N_m^{(r)}} \sum_{i: \alpha^{(r)}(X_i) = m} \rho_0(X_i, g),$$

where $N_m^{(r)} \triangleq |\{i : \alpha^{(r)}(X_i) = m\}|$, i.e., set $\beta^{(r)}(m) = g_m^{(r)} = \arg \min_{g \in \mathcal{G}} E[\rho_0(X, g) | \alpha^{(r)}(X) = m]$.

- 4) **Optimal length function:** if $N_m^{(r)} > 0$, let $\ell^{(r)}(m) = -\ln p_m^{(r)}$, where $p_m^{(r)} = N_m^{(r)} / N$. If $N_m^{(r)} = 0$, remove the corresponding cell from the code and decrease $|\mathcal{M}|$ by 1.
- 5) **Test:** compute the average ρ -distortion D_r with the code $(\alpha^{(r)}, \beta^{(r)}, \ell^{(r)})$. If $(D_{r-1} - D_r) / D_{r-1} < \epsilon$, quit. Otherwise, go to Step 2 and continue.

With a judicious choice of the initial codebook and length function, this algorithm yields a finite mixture model $\{(g_m, p_m) : m \in \mathcal{M}\}$ as a good “fit” to the empirical distribution of the data in the sense of near-optimal trade-off between the local mismatch and complexity.

IV. APPLICATION TO NLDR

Given a training sample $X^N = (X_1, \dots, X_N)$ of “raw” n -dimensional data and assuming its intrinsic dimension $k < n$ is known, our goal is to determine two mappings, $v : \mathbb{R}^n \rightarrow \mathbb{R}^k$ and $w : \mathbb{R}^k \rightarrow \mathbb{R}^n$, where v maps high-dimensional vectors to their dimension-reduced versions and w maps back to the high-dimensional space. In general, the dimension-reducing map v

entails loss of information, so $w(v(x)) \neq x$. Therefore we will be interested in the average distortion incurred by our scheme, $\bar{d}(v, w) \triangleq \mathbb{E}[d(X, w(v(X)))]$, where $d : \mathbb{R}^n \times \mathbb{R}^n \rightarrow [0, \infty)$ is a suitable distortion measure on pairs of n -vectors, e.g., the squared Euclidean distance, and the expectation is w.r.t. the empirical distribution of the sample.

A. Mixture model of a stochastic embedding

The first step is to use the above quantization scheme to fit a complexity-regularized Gaussian mixture model to the training sample. Our class \mathcal{G} of admissible model pdf's will be the set of all n -dimensional Gaussians with nonsingular covariance matrices, $\mathcal{G} = \{\mathcal{N}(x; \mu, K) : \mu \in \mathbb{R}^n, \det K > 0\}$, and for each finite set $\Gamma \subset \mathcal{G}$ we shall define a regularization functional $\Phi_\Gamma : \Gamma \rightarrow [0, \infty)$ that penalizes those $g \in \Gamma$ that are “geometrically complex” relative to the rest of Γ .

The idea of “geometric complexity” can be motivated [5], [7] by the example of the Gaussian mixture model from Sect. II. The covariance matrix of the l th component, $A_l A_l^t + \Sigma_l$, is invariant under the mapping $A_l \mapsto A_l R$, where R is a $k \times k$ orthogonal matrix, i.e., $RR^t = I$. In geometric terms, a copy of the orthogonal group O_k associated with the l th component of the mixture is the group of rotations and reflections in the tangent space to M at $\varphi_l^{-1}(0)$. Thus, the log-likelihood term in ρ_0 is not affected by assigning arbitrary and independent orientations to the tangent spaces associated with the components of the mixture. However, since our goal is to model the intrinsic *global* geometry of the data, it should be possible to smoothly glue together the local data provided by our model. We therefore require that the orientations of the tangent spaces at “nearby” points change smoothly as well. (In fact, one has to impose certain continuity requirements on the orientation of the tangent spaces in order to define measure and integration on the manifold [4, Ch. XI].)

Given a finite set $\Gamma \subset \mathcal{G}$, we shall define the regularization functional $\Phi_\Gamma : \Gamma \rightarrow [0, \infty)$ as

$$\Phi_\Gamma(g) \triangleq \sum_{g' \in \Gamma \setminus \{g\}} \kappa(\mu_g, \mu_{g'}) D(g' \| g), \quad (2)$$

where $\kappa : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^+$ is a smooth positive symmetric kernel such that $\kappa(x, x') \rightarrow 0$ as $\|x - x'\| \rightarrow \infty$, and

$$D(g' \| g) = \frac{1}{2} (\ln \det(K_g^{-1} K_{g'}) + \text{Tr}(K_g^{-1} K_{g'}) + (\mu_g - \mu_{g'})^t K_g^{-1} (\mu_g - \mu_{g'}) - n)$$

is the relative entropy between two Gaussians. Possible choices for the kernel κ are the inverse Euclidean distance $\kappa(x, x') = \|x - x'\|^{-1}$ [8], a Gaussian kernel $\kappa(x, x') = \mathcal{N}(x - x'; 0, \sigma^2 I)$ for a suitable value of σ [7], [8] or a compactly supported “bump” $\kappa(x, x') = \psi_{r_1, r_2}(x - x')$, where ψ_{r_1, r_2} is an infinitely differentiable reflection-symmetric function that is identically zero everywhere outside a closed ball of radius r_2 and one everywhere inside an open ball of radius $r_1 < r_2$. The relative entropy serves as a measure of position and orientation alignment of the tangent spaces, while the smoothing kernel ensures that more weight is assigned to “nearby” components.

This complexity functional is a generalization of the “global coordination” prior of Brand [7] to mixtures with unequal component weights.

With these definitions of \mathcal{G} and Φ_Γ , the ρ_0 -distortion for a codebook $\Gamma = \{g_m : m \in \mathcal{M}\}$ and a length function ℓ is

$$\rho_0(x, m) = \frac{1}{2} \ln \det K_m + \frac{1}{2} (x - \mu_m)^t K_m^{-1} (x - \mu_m) + \ell(m) + \sum_{m' \in \mathcal{M} \setminus \{m\}} \kappa(\mu_m, \mu_{m'}) D(g_{m'} \| g_m),$$

where we have also removed the $(n/2) \ln(2\pi)$ term as it does not affect the encoder. The effect of the geometric complexity term is to curve the boundaries of the partition cells according to locally interpolated “nonlocal information” about the rest of the codebook. Determining the Lloyd centroids for the decoder will involve solving $|\mathcal{M}|$ simultaneous nonlinear equations for the means and the same number of equations for the covariance matrices. For computational efficiency we can use the kernel data from the previous iteration, which would sacrifice optimality but avoid nonlinear equations.

B. Design of reduction and reconstruction maps

The output of the previous step is a Gauss mixture model $\{(g_m, p_m) : m \in \mathcal{M}\}$ and a partition $\mathcal{R} = \{R_m\}$ of \mathbb{R}^n . Suppose that for each $m \in \mathcal{M}$ the eigenvectors $e_1^{(m)}, \dots, e_n^{(m)}$ of K_m are numbered in the order of decreasing eigenvalues, $\lambda_1^{(m)} \geq \dots \geq \lambda_n^{(m)}$. The next step is to design the dimension-reducing map v and the reconstruction map w . One method, proposed by Brand [7], is to use the mixture model of the underlying pdf [obtained in his case by an EM algorithm with a prior corresponding to the average of the complexity $\Phi_\Gamma(g)$ over the entire codebook and with equiprobable components of the mixture] to construct a mixture of local affine transforms, preceded by local Karhunen-Loève transforms, as a solution to a weighted least-squares problem.

However, we can use the encoder partition \mathcal{R} directly: for each $m \in \mathcal{M}$, let $v_m(x) \triangleq \Pi_m(x - \mu_m)$, where Π_m is the projection onto the first k eigenvectors of K_m , and then define $v(x) = \sum_{m \in \mathcal{M}} 1_{\{x \in R_m\}} v_m(x)$. This approach is similar to local principal component analysis of Kambhatla and Leen [9], except that their quantizer was not complexity-regularized and therefore the shape of the resulting Voronoi regions was determined only by local statistical data. We can describe the operation of dimension reduction (feature extraction) as an encoder $\hat{v} : \mathbb{R}^n \rightarrow \mathcal{M} \times \mathbb{R}^k$, so that $\hat{v}(x) = (m(x), v_{m(x)}(x))$, where α is the minimum-distortion encoder for the ρ_0 -distortion.

The corresponding reconstruction operation can be designed as a decoder $\hat{w} : \mathcal{M} \times \mathbb{R}^k \rightarrow \mathbb{R}^n$ which receives a pair (m, u) , $m \in \mathcal{M}$, $u \in \mathbb{R}^k$, and computes $w_m(u) = \mu_m + \sum_{i=1}^k \langle u, e_i^{(m)} \rangle e_i^{(m)}$, where $\langle \cdot, \cdot \rangle$ denotes the usual scalar product in \mathbb{R}^k .

This encoder-decoder pair is a composite Karhunen-Loève transform coder matched to the mixture source $g = \sum_m p_m g_m$. If the data alphabet \mathcal{X} is compact, then the squared-error distortion is bounded by some $A > 0$, and

the mismatch due to using this composite coder on the disjoint mixture source $f = \sum_m p_m f_m$ can be bounded from above by $A\|f - g\|_1$, where $\|\cdot\|_1$ is the L_1 norm. Provided that the mixture g is optimal for f in the sense of minimizing the ρ -distortion, we can use Pinsker's inequality [10, Ch. 5] $\|f - g\|_1 \leq \sqrt{2D(f\|g)}$ and convexity of the relative entropy to further bound the mismatch by $A\sqrt{2(\bar{I}_f(\mathcal{R}, \Gamma) - \mu \sum_m p_m \Phi_\Gamma(g_m))}$.

Note that the maps v and w are not smooth, unlike the analogous maps of Brand [7], [8]. This is an artifact of the hard partitioning used in our scheme. However, hard partitioning has certain advantages: it allows for use of composite codes [6] and nonlinear interpolative vector quantization [11] if additional compression of dimension-reduced data is required. Moreover, the lack of smoothness is not a problem in our case because we can use kernel interpolation techniques to model the geometry of dimension-reduced data by a smooth manifold, as explained next.

C. Manifold structure of dimension-reduced data

Our use of mixture models has been motivated by certain assumptions about the structure of stochastic embeddings of low-dimensional manifolds into high-dimensional spaces. In particular, given an n -dimensional Gaussian mixture model $\{(g_m, p_m) : m \in \mathcal{M}\}$, we can associate to each component of the mixture a chart of the underlying manifold, such that the image of the chart in \mathbb{R}^k is an open ball of radius $r_m = (\lambda_1^{(m)})^{1/2}$ centered at the origin, and we can take the first k eigenvectors of the covariance matrix of g_m as coordinate axes in the tangent space to the manifold at the inverse image of $0 \in \mathbb{R}^k$ under the m th chart. Owing to geometric complexity regularization, the orientations of tangent spaces change smoothly as a function of position.

Ideally, one would like to construct a smooth manifold consistent with the given descriptions of charts and tangent spaces. However, this is a fairly difficult task since we not only have to define a smooth coordinate map φ_m for each chart, but also make sure that these maps satisfy the chart compatibility condition. Instead, we can construct the manifold *implicitly* by gluing the coordinate frames of the tangent spaces into an object having a smooth inner product.

Specifically, let us fix a sufficiently small $\delta > 0$, and let ψ_m be an infinitely differentiable function that is identically zero everywhere outside a closed ball of radius r_m and one everywhere inside an open ball of radius $r_m - \delta$, with both balls centered at $\Pi_m \mu_m$. Let $\eta_m(u) \triangleq \frac{p_m \psi_m(u)}{\sum_{m \in \mathcal{M}} p_m \psi_m(u)}$. The inner product of two vectors $u, u' \in \mathbb{R}^k$, treated as elements of the tangent space $T_{\varphi_m^{-1}(0)}M$, is given by $\langle u, u' \rangle_m = \sum_{i=1}^k \langle u, e_i^{(m)} \rangle \langle e_i^{(m)}, u' \rangle$. Then for each $y \in \mathbb{R}^k$ the map $g_y : \mathbb{R}^k \times \mathbb{R}^k \rightarrow [0, \infty)$,

$$g_y(u, u') \triangleq \sum_{m \in \mathcal{M}} \eta_m(y + \Pi_m \mu_m) \langle u, u' \rangle_m,$$

is a symmetric form, which is positive definite whenever $\eta_m(y + \Pi_m \mu_m) \neq 0$ for at least one value of m . In addition,

the map $y \mapsto g_y(\cdot, \cdot)$ is smooth. In this way, we have implicitly defined a *Riemannian metric* [4, Ch. VII] on the underlying manifold. The functions η_m form a so-called *smooth partition of unity*, which is the only known way of gluing together local geometric data to form smooth objects [4, Ch. II].

In geometric terms, $\eta_m(y + \Pi_m \mu_m) = 0$ for all m if and only if $y \in \mathbb{R}^k$ is an image under the dimension-reduction map of a point in \mathbb{R}^n whose first k principal components w.r.t. each Gaussian in the mixture model fall outside the covariance ellipsoid of that Gaussian. If the mixture model is close to optimum, this will happen with negligible probability. A practical advantage of this feature of our scheme is in rendering it robust to outliers.

V. CONSISTENCY AND CODEBOOK DESIGN

Our mixture modeling scheme can also be used to estimate the “true” but unknown pdf f^* of the high-dimensional data, if we assume that f^* belongs to some fixed class \mathcal{F} . Indeed, the empirically designed codebook $\Gamma = \{g_m : m \in \mathcal{M}\}$ of Gaussian pdf's, the corresponding component weights $\{p_m\}$, and the mixture $g = \sum_{m \in \mathcal{M}} p_m g_m$ are random variables since they depend on the training sample X^N . We are interested in the quality of approximation of f^* by the mixture $g \equiv g(X^N)$.

Following Moulin and Liu [12], we use the relative-entropy loss function $D(f^*\|g)$. We shall give an upper bound on the loss in terms of the *index of resolvability* [12]

$$R_{\mu, N}(f^*) \triangleq \min_{m \in \mathcal{M}} \left[D(f^*\|g_m) + \frac{\mu L(g_m)}{N} \right],$$

where $L(g_m) \triangleq \Phi_\Gamma(g_m) - \ln p_m$, which quantifies how well f^* can be approximated, in the relative-entropy sense (and, by Pinsker's inequality, in L_1 sense), by a Gaussian of moderate geometric complexity relative to the rest of the codebook. We have the following result:

Theorem V.1: Let the codebook $\Gamma = \{g_m : m \in \mathcal{M}\}$ of Gaussian pdf's be such that the log-likelihood ratios $U_m \triangleq -\ln(f^*(X)/g_m(X))$ uniformly satisfy the *Bernstein moment condition* [10], i.e., there exists some $h > 0$ such that $\mathbb{E}|U_m - \mathbb{E}U_m|^k \leq (1/2) \text{Var}(U_m) k! h^{k-2}$ for all $k \geq 2$. Let $M(f^*)$ be the smallest number such that $\text{Var}(U_m) \leq -M(f^*) \mathbb{E}U_m$ for all $m \in \mathcal{M}$ (owing to the Bernstein condition, it is nonnegative and finite). Then, for any $\mu > h + M(f^*)/2$ and $\delta > 0$,

$$\Pr \left\{ D(f^*\|g) \leq \frac{1 + \alpha}{1 - \alpha} R_{\mu, N}(f^*) + \frac{2\mu \ln \frac{|\mathcal{M}|}{\delta}}{(1 - \alpha)N} \right\} \geq 1 - 2\delta, \quad (3)$$

where $\alpha = \frac{M(f^*)}{2(\mu - h)}$. The expected loss satisfies

$$\mathbb{E}[D(f^*\|g)] \leq \frac{1 + \alpha}{1 - \alpha} R_{\mu, N}(f^*) + \frac{4|\mathcal{M}|\mu}{(1 - \alpha)N}. \quad (4)$$

The probabilities and expectations are all w.r.t. the pdf f^* .

Proof: Due to the fact that $\Phi_\Gamma(g_m) \geq 0$ for all $m \in \mathcal{M}$, the composite complexity $L(g_m)$ satisfies the Kraft inequality. Then we can use a strategy similar to that of Moulin and Liu

[12] to prove that

$$\Pr \left\{ D(f^* \| g_m) \geq \frac{1+\alpha}{1-\alpha} R_{\mu,N}(f^*) + \frac{2\mu \ln \frac{|\mathcal{M}|}{\delta}}{(1-\alpha)N} \right\} \leq \frac{2\delta}{|\mathcal{M}|}$$

for each $m \in \mathcal{M}$. Hence, by the union bound

$$D(f^* \| g_m) \leq \frac{1+\alpha}{1-\alpha} R_{\mu,N}(f^*) + \frac{2\mu \ln \frac{|\mathcal{M}|}{\delta}}{(1-\alpha)N}$$

for all $m \in \mathcal{M}$, except for an event of probability at most 2δ . By convexity of the relative entropy, $D(f^* \| g_m) \leq C$ for all $m \in \mathcal{M}$ implies that $D(f^* \| g) \leq C$ for $g = \sum_{m \in \mathcal{M}} p_m g_m$. Therefore

$$D(f^* \| g) \leq \frac{1+\alpha}{1-\alpha} R_{\mu,N}(f^*) + \frac{2\mu \ln \frac{|\mathcal{M}|}{\delta}}{(1-\alpha)N}$$

with probability at least $1 - 2\delta$. To prove (3), we use the fact [10] that if Z is a random variable with $E|Z| < \infty$, then $E[Z] \leq \int_0^\infty \Pr[Z \geq t] dt$. We let $Z = D(f^* \| g) - \frac{1+\alpha}{1-\alpha} R_{\mu,N}(f^*)$ and choose $\delta = |\mathcal{M}| e^{-\frac{Nt(1-\alpha)}{2\mu}}$. Then $E[Z] \leq \frac{4|\mathcal{M}|\mu}{(1-\alpha)N}$, which proves (4). ■

To discuss consistency in the large-sample limit, consider a sequence of empirically designed mixture models $\{(g_m^{(N)}, p_m^{(N)}) : m \in \mathcal{M}^{(N)}\}$. This is different from the usual empirical quantizer design, where we increase the training set size but keep the number of quantizer levels fixed. The scheme is consistent in the relative-entropy sense if $E D(f^* \| g^{(N)}) \rightarrow 0$ as $N \rightarrow \infty$, where $g^{(N)} = \sum_{m \in \mathcal{M}^{(N)}} p_m^{(N)} g_m^{(N)}$ and the expectation is with respect to f^* .

A sufficient condition for consistency can be determined by inspection of the upper bound in Eq. (4). Specifically, we require that the codebooks $\Gamma^{(N)}$ satisfy: (a) $\max_{m \in \mathcal{M}^{(N)}} L(g_m^{(N)}) = o(N)$, (b) $\min_{m \in \mathcal{M}^{(N)}} D(f^* \| g_m) = o(1)$ for all $f^* \in \mathcal{F}$, and (c) $|\mathcal{M}^{(N)}| = o(N)$. Condition (c) can be satisfied by initializing the Lloyd algorithm by a codebook of size much smaller than the training set size N , which is usually done in practice in order to ensure good training performance. The first two conditions can also be easily met in many practical settings.

Consider, for instance, the class \mathcal{F} of all pdf's supported on a compact $\mathcal{X} \subset \mathbb{R}^n$ and Lipschitz-continuous with Lipschitz constant c . Then, if we take as our class of admissible Gaussians $\mathcal{G} = \{\mathcal{N}(x; \mu, K) : \mu \in \mathcal{X}, c_1 \leq \det K \leq c_2\}$ for suitably chosen constants $c_1, c_2 > 0$ independent of N , the relative entropy $D(g \| g')$ of any two $g, g' \in \mathcal{G}$ can be bounded independently of N , and condition (a) will be met with proper choice of the component weights. Condition (b) is likewise easy to meet since the maximum value of any $f^* \in \mathcal{F}$ depends only on the set \mathcal{X} , the Lipschitz constant c , and the dimension n .

In general, the issue of optimal codebook design is closely related to the problem of universal vector quantization [13]: we can consider, e.g., a class \mathcal{F} of pdf's with disjoint supports contained in a compact $\mathcal{X} \subset \mathbb{R}^n$. Then a sequence of Gaussian codebooks that yields a consistent estimate of each $f^* \in \mathcal{F}$ in the large-sample limit is weakly minimax universal [13] for

\mathcal{F} and can also be used to quantize any source contained in the L_1 -closed convex hull of \mathcal{F} .

VI. DISCUSSION

We have introduced a complexity-regularized quantization approach to NLDR. One advantage of this scheme over existing methods for NLDR based on Gaussian mixtures, e.g., [7], is that, instead of fitting a Gauss mixture to the entire sample, we design a codebook of Gaussians that provides a good trade-off between local adaptation to the data and global geometric coherence, which is key to robust geometric modeling. Complexity regularization is based on a kernel smoothing technique that allows for a meaningful geometric description of dimension-reduced data by means of a Riemannian metric and is also robust to outliers. Moreover, to our knowledge, the consistency proof presented here is the first theoretical asymptotic consistency result applied to NLDR.

Work is currently underway to implement the proposed scheme for applications to image processing and computer vision. Also planned is future work on a quantization-based approach to estimating the intrinsic dimension of the data and on assessing asymptotic *geometric* consistency of our scheme in terms of the Gromov-Hausdorff distance between compact metric spaces [14].

Acknowledgment. I would like to thank Svetlana Lazebnik and Prof. Pierre Moulin for useful discussions. This research has been supported by the Beckman Postdoctoral Fellowship.

REFERENCES

- [1] J. Tenenbaum, V. de Silva, and J. Langford, "A global geometric framework for nonlinear dimensionality reduction," *Science*, vol. 290, pp. 2319–2323, December 2000.
- [2] S. Roweis and L. Saul, "Nonlinear dimensionality reduction by locally linear embedding," *Science*, vol. 290, pp. 2323–2326, December 2000.
- [3] E. Levina and P. Bickel, "Maximum likelihood estimation of intrinsic dimension," in *Adv. Neural Inform. Processing Systems*, L. Saul, Y. Weiss, and L. Bottou, Eds., vol. 17. Cambridge, MA: MIT Press, 2005.
- [4] S. Lang, *Differential and Riemannian Manifolds*. New York: Springer-Verlag, 1995.
- [5] S. Roweis, L. Saul, and G. Hinton, "Global coordination of locally linear models," in *Adv. Neural Inform. Processing Systems*, T. Dietterich, S. Becker, and Z. Ghahramani, Eds., vol. 14. Cambridge, MA: MIT Press, 2002, pp. 889–896.
- [6] R. Gray and T. Linder, "Mismatch in high-rate entropy-constrained vector quantization," *IEEE Trans. Inform. Theory*, vol. 49, no. 5, pp. 1204–1217, May 2003.
- [7] M. Brand, "Charting a manifold," in *Adv. Neural Inform. Processing Systems*, S. Becker, S. Thrun, and K. Obermayer, Eds., vol. 15. Cambridge, MA: MIT Press, 2003, pp. 977–984.
- [8] —, "Continuous nonlinear dimensionality reduction by kernel eigenmaps," in *Int. Joint Conf. Artif. Intel.*, 2003.
- [9] N. Kambhatla and T. Leen, "Dimension reduction by local principal component analysis," *Neural Comput.*, vol. 9, pp. 1493–1516, 1997.
- [10] L. Devroye and G. Lugosi, *Combinatorial Methods in Density Estimation*. New York: Springer-Verlag, 2001.
- [11] A. Gersho, "Optimal nonlinear interpolative vector quantization," *IEEE Trans. Commun.*, vol. 38, no. 9, pp. 1285–1287, September 1990.
- [12] P. Moulin and J. Liu, "Statistical imaging and complexity regularization," *IEEE Trans. Inform. Theory*, vol. 46, no. 5, pp. 1762–1777, August 2000.
- [13] P. Chou, M. Effros, and R. Gray, "A vector quantization approach to universal noiseless coding and quantization," *IEEE Trans. Inform. Theory*, vol. 42, no. 4, pp. 1109–1138, July 1996.
- [14] P. Petersen, "Gromov-Hausdorff convergence of metric spaces," in *Summer Inst. Diff. Geom.*, ser. Proc. Symposia Pure Math. Amer. Math. Soc., 1990, pp. 489–505.